

continue “heroic” measures like a ventilator, “ordinary” treatments like antibiotics and intravenous fluids should be continued.*

Turning off the ventilator is murder only if it violates autonomy. It might seem that complying with Mr. C’s reasoned and well-informed request could hardly violate his autonomy, but as before, his death is a violation of agency because it is inconsistent with any and every action plan he might adopt. It makes no difference whether his death results from “active” or “passive” euthanasia, or whether it is self-inflicted (suicide). All destroy agency. The distinction of “heroic” and “ordinary” efforts is part of popular culture, but it, too, has no bearing on the ethical analysis. Withdrawing any effort that would save his life destroys agency.

One can avoid violation of Mr. C’s autonomy only if he has literally reached the point where he has no action plans other than his own demise. He is forming no intentions, large or small, for tomorrow or the next moment, other than persuading his caretakers to let him die. I am not certain this is psychologically possible, or even a logically coherent notion; however, if it is, and if Mr. C finds himself in this state of mind, then he can opt for his own death without violating autonomy. If others can ascertain that he is in this state of mind, they can bring about his death without violating autonomy. There may be other considerations, such as whether his continued existence would benefit others in some way, but absent these, suicide and euthanasia are ethical in this case.

ALLOCATING MEDICAL RESOURCES

A pervasive problem in medicine is the allocation of scarce resources. It arises in a number of contexts, ranging from triage to organ transplantation to public funding. Healthcare is an area of mind-boggling complexity, and I cannot pretend to supply an adequate analysis of just allocation in these pages. However, I can get the discussion going, and I will do so by focusing on triage. This will illustrate how ethical arguments can be brought to bear in this area. It will also provide an opportunity to address an issue of distributive justice.

Triage is usually understood as the allocation of resources in an emergency situation, such as on a battlefield (where the concept apparently

* Based on Case 15.1 of B. Lo, cited earlier.

originated) or in a hospital emergency room. Traditionally, it divides patients into three categories: those who will survive without treatment, those who will perish even with treatment, and those who might be saved by treatment. The last category of patients receives first priority.

Modern triage systems typically recognize several categories, perhaps determined by a *trauma score*. A common system is color coded: black for patients who are deceased or beyond help, red for those who require immediate life-saving intervention, yellow for patients who are stable but must be continuously watched and perhaps re-triaged, green for patients who can be returned to later, and white for patients who can be released from care. The red category receives the highest priority, followed by yellow, green, white, and black. The classification can be refined in various ways. Patients in the black category may receive limited attention, such as an effort to resuscitate those without vital signs, or palliative care to those still alive, perhaps in the form of pain medication. Patients with severed fingers may be given priority similar to that of the red category, even though their injury is not life threatening, so as to restore the fingers while this is still possible.

These triage policies seem roughly to apply a utilitarian rule: Adopt the treatment protocol that results in the greatest overall benefit to patients, given scarce resources. The black category receives little or no priority, because diverting resources to those who would benefit little from them reduces net utility. The yellow category receives priority over the green category, because lavishing attention on those who don't need it now benefits them only marginally while potentially harming those who need close monitoring, and so forth.

Utility should indeed be a central consideration in triage, but respect for autonomy must also be considered. Unnecessarily neglecting a patient who will die without treatment violates the Joint Autonomy Principle, because it is inconsistent with any action plan the patient might have. Or more precisely, any patient who doctors are rationally constrained to believe will die without treatment must be given highest priority. This is regardless of the utilitarian cost. If an entire medical team is required to save the one patient who has a life-threatening injury, they must grit their teeth and focus on this one patient—even if it means neglecting all other patients, and even if these patients will suffer serious (but nonlethal) consequences from delayed treatment. I am not sure the Utilitarian Principle calls for this, but the Joint Autonomy Principle demands it. It also seems consistent with accepted triage policies.

Autonomy principles seem to have less to say in other scenarios. While delaying treatment in a life-threatening case violates autonomy, this is not clearly true when the injury is nonlethal. Delayed treatment may be consistent with the limited range of action plans that are possible for an injured patient. The patient might still be said to have an action plan of getting married next month or retiring next year, because it may be rational to believe that these are possible despite the injury. Yet these plans may be consistent with delayed treatment. At any rate, the medics are not rationally constrained to believe that delayed treatment is inconsistent with the patient's plans, if only because in an emergency situation, they have no specific knowledge of any patient's action plans. They can therefore revert to utilitarian criteria without fear of violating autonomy.*

We should probably make an exception to this conclusion when delayed treatment would result in permanent impairment. The medics may be rationally constrained to believe that such an impairment would be inconsistent with the patient's current action plans, which makes delay a violation of autonomy. In such cases, the patient should receive priority similar to that of those in the red category. Interestingly, this seems to justify the practice of giving priority to cases in which a limb must be quickly reattached to be saved.

Autonomy principles provide no guidance when there are too many life-threatening cases to treat. If the medics are certain that the seriously injured patients will die without attention, then the dilemma of which ones to treat is similar to the trolley car dilemmas discussed in Chapter 6. Any selection for treatment violates autonomy, because it is inconsistent with any possible action plan of the patients who are not selected. This doesn't say that the medics are somehow at fault. It only says that no autonomous choice is possible, and the medics can only follow instinct or prior training. Fortunately, the situation is often less stark. The medics

* Interestingly, the situation may change in the unusual situation in which the medics are acquainted with the patients. They may know that some patients have action plans that would be thwarted by delayed treatment, and others do not. The Joint Autonomy Principle requires them to give priority to the first group. A more difficult dilemma arises when the medics know that some patients have such action plans but are not acquainted with the other patients and are therefore not rationally constrained to believe that they have any such plans. Then the medics must give priority to the patients they know. This strikes us as unfair because it plays favorites, and yet it seems to follow from the Joint Autonomy Principle. We must be prepared to accept ethical conclusions we don't like. At any rate, none of these exceptional situations need be addressed in a triage policy, because acquaintance with patients is extremely rare in an emergency situation.

may only be able to judge that certain patients are *likely* to die without immediate treatment, in which case there is no violation of autonomy. The medics can follow a utilitarian policy.

Formulating a utilitarian policy, meanwhile, is not as straightforward as it may seem. It is not enough simply to give higher priority to patients with more serious injuries, because the likely benefit may be smaller. The proper criterion is to maximize expected utility, which is the product of the utility gained by prompt treatment and the probability of achieving this gain, summed over all patients.

Measuring utility can also be tricky. Prompt treatment of one patient may extend life many years but with poor quality of life, whereas prompt treatment of another patient may yield fewer additional years but with good health. Even if future quality of life is not ordinarily considered in triage situations, it can be a major factor when formulating health policy, as for example when a national health agency allocates funding. A common approach is to estimate the number of additional *quality-adjusted life years* (QALYs) yielded by treatment.* This number is obtained by assigning a quality-of-life assessment to each year the patient is expected to live, first assuming immediate treatment, and then assuming no immediate treatment. The assessments vary from 0 (unconscious) to 1 (good health). There are various methods for quantifying quality of life, based partly on objective criteria and partly on patient self-assessment.† The quality-of-life assessments are then summed over the expected years of life under either

* J. Broome, Fairness and QALYs, *Philosophy and Medical Welfare* 3 (1988): 57–73; P. Dolan, The measurement of individual utility and social welfare, *Journal of Health Economics* 17 (1998): 39–52.

† Two early efforts at assessing quality of life in a medical contest are described in A. Williams, Economics of coronary bypass grafting, *British Medical Journal* 291 (1985): 326–329; R. W. Evans, D. L. Manninen, L. P. Garrison, L. G. Hart, C. R. Blagg, R. A. Gutman, A. R. Hill, and E. G. Lowrie, The quality of life of patients with end-stage renal disease, *New England Journal of Medicine* 312 (1985): 553–559. Recently used methods for assessing quality of life include time tradeoff, described in K. Burström; M. Johannesson and F. Diderichsen, A comparison of individual and social time-trade-off values for health states in the general population, *Health Policy* 76 (2006): 359–370; a lottery comparison, described in G. W. Torrance, Measurement of health state utilities for economic appraisal, *Journal of Health Economics* 5 (1986): 1–30; the EQ-5D Questionnaire, described in EuroQol Group, EuroQol: A new facility for the measurement of health-related quality of life, *Health Policy* 16 (1990): 199–208; and a simple subjective scale similar to that used when patients are asked to report pain level. Other measures are referenced in D. H. Feeny, E. Eckstrom, and E. P. Whitlock, Patient-reported outcomes, health-related quality of life, and function: An overview of measurement properties, a chapter in *A Primer for Systematic Reviewers on the Measurement of Functional Status and Health-Related Quality of Life in Older Adults*, Agency for Healthcare Research and Quality, available online, 2013.

scenario (with and without treatment). The benefit of immediate treatment is the difference in total QALYs between the two scenarios.

Healthcare policy can also raise a justice issue. Suppose a few patients with a serious disease can probably be cured, but at great expense. The same funds could create greater total utility by financing a cure for a mild form of sniffles that afflicts millions. Since autonomy is not at stake, the Utilitarian Principle instructs us to cure the sniffles rather than treat the serious disease. This may not seem to be a just distribution of resources. Similar issues arise constantly in social policy discussions. For example, a government might boost economic growth and create more total utility by abolishing the minimum wage and allowing an impoverished underclass to toil for the benefit of others. One might question whether a minimum wage actually reduces total utility, but the point here is that abolishing it may be unjust even if it is utilitarian.

This suggests that the Utilitarian Principle should be restricted to considering action plans that result in just distributions. We have already restricted it to actions that are generalizable and respect autonomy, and perhaps a third condition is necessary. But exactly what is a just distribution?

By far the best-known criterion of distributive justice is the Difference Principle of John Rawls.^{*} It states, very roughly, that differences in welfare are justified only to the extent they are necessary to maximize the welfare of the worst-off. Equality of outcomes would be nice, but perhaps a society needs incentives for people to work hard, or else everyone will be impoverished. This creates inequality, because not everyone is willing or able to earn success. The Difference Principle tells us that it is okay for some people to be less advantaged than others, but only to a degree that lifts the bottom as high as possible.

Rawls advances a social contract argument for this principle. In a nutshell, it says that a rational social policy must be one on which rational agents can agree up front, in an *original position*, without knowing how it will affect them personally. It is as though everyone is attending a constitutional convention to arrive at a social contract. They negotiate behind

^{*} J. Rawls, *A Theory of Justice*, Cambridge, MA, Harvard University Press, 1971; *Political Liberalism*, New York, Columbia University Press, 1993. For ethical analyses of inequality, see L. S. Tempkin, *Inequality*, Oxford, Oxford University Press, 1993; H. G. Frankfurt, *On Inequality*, Princeton, NJ, Princeton University Press, 2015.

a *veil of ignorance* (as Rawls puts it) as to who will occupy each position in the new society. Any delegate to the convention could end up a corporate CEO or a migrant farm worker and must be willing to endorse the contract in either case. Rawls argues that such an endorsement can be rational only if the farm worker at the bottom of the economy would have been even worse off in any other social arrangement. Thus, social policy should maximize the welfare of the worst-off. Perhaps a wealth redistribution program would disincentivize workers, reduce overall economic output, and move the bottom rung of the ladder even lower. If so, the social contract satisfies the Difference Principle despite its structural inequality.

The Difference Principle instructs us to treat the patient with a serious disease rather than cure the sniffles. This maximizes the welfare of the patient who is worse off. It fails to maximize the welfare of patients with the sniffles, but they don't count, because they are better off than the patient with the serious disease.

The Difference Principle reflects the spirit of this book by insisting that the rationality of a choice should not depend on who one is. Yet the principle must be carefully formulated to withstand the intense scrutiny it has received, and it remains controversial. Whatever the outcome of this debate, there is a more basic problem for our project: The Difference Principle doesn't so much limit the Utilitarian Principle as replace it, which is problematic because the Utilitarian Principle has strong arguments in its favor.

The problem can be illustrated by a second healthcare dilemma. Suppose a few persons have a certain rare, incurable form of cancer, and their lives can be extended a week or two by giving them an extremely costly treatment. Or suppose we could use the same resources to develop a vaccine that won't save any lives but will relieve millions from a painful tropical disease that results in significant lost productivity. The Difference Principle requires that we spend the money on the cancer patients, because they are worse off, yet the vaccine results in far greater utility and seems the rational choice.*

* This situation is even worse than this. The same logic that leads to the "maximin" criterion of the Difference Principle (maximize the minimum) can also justify a "lexicographic maximum." This means that one should maximize the welfare of the worst off, hold it fixed while maximizing the welfare of the second worst off, and so on through all economic levels. This leaves even less room for utilitarian considerations to operate.

The ideal would be to find some larger principle that reconciles the Utilitarian and Difference Principles and that can be grounded in the logical structure of action.* This is mainly an issue for social policy rather than the type of individual decision-making I address in this book, but it poses a fundamental research question for normative ethics.

* One way to forge a compromise between the Rawlsian and utilitarian criteria is to maximize a weighted average of the minimum and total utility. It is highly unclear, however, what kind of weighting is appropriate. A recent proposal is to count the utility of everyone whose utility is within D of the lowest utility as having the lowest utility and then maximize total utility. When D is zero, we have a purely utilitarian criterion, and when D is infinite, we have a purely Rawlsian criterion. A specialization of the idea was first proposed for the case of two individuals by A. Williams and R. Cookson, Equity in health, in A. J. Culyer and J. P. Newhouse (Eds.), *Handbook of Health Economics*, Philadelphia, PA, Elsevier Science, 2000. It was generalized to larger populations and applied to healthcare provision by J. N. Hooker and H. P. Williams, Combining equity and utilitarianism in a mathematical programming model, *Management Science* 58 (2012): 1682–1693.